

**AbNET,  
a fault-tolerant fiber optic communication system**

**Harold Kirkham**  
**Jet Propulsion Laboratory**  
**California Institute of Technology**  
**Pasadena, California, USA**  
**Phone**        **+ 1 8183549699**  
**Fax**         **- 1 8183934820**  
**e-mail**       **Harold.Kirkham@jpl.nasa.gov**

**Eddie Hsu**  
**Jet Propulsion Laboratory**  
**California Institute of Technology**  
**Pasadena, California, USA**  
**Phone**        **+ 1 8183543212**  
**Fax**         **+ 1 8183934820**  
**e-mail**       **Hsu@jplpto.jpl.nasa.gov**

# **AbNET, a fault-tolerant fiber optic communication system**

**Harold Kirkham  
Jet Propulsion Laboratory  
California Institute of Technology  
Pasadena, California, USA**

**Eddie Hsu  
Jet Propulsion Laboratory  
California Institute of Technology  
Pasadena, California, USA**

**ABSTRACT** The design of a fiber optic communication network originally intended for monitoring and control in power distribution systems is discussed. By appropriate choice of protocols, a fault-tolerant system can be built that operates in any arbitrary network configuration.

The network, called AbNET, is a packet-based distributed protocol system. Flooding is used for maximum failure tolerance. Hierarchical (master-slave) polling controls access to the system. This supports many data acquisition and control applications. The protocols allow multiple adjacent masters to share resources. A service reports the network's configuration to the master, where changes can be logged, and action taken if needed.

The system is transparent to the user, and maintains no record of the clients it serves. It is fast enough for many industrial control applications. Because hierarchical access control is used, peer-to-peer communications, such as for MMS systems, do not map well into the protocols, but can be accomplished via software at the master station.

Only a small number of fiber cables is needed for a high reliability system. In many industrial applications, where the inter-node distance is not large, fiber is unlikely to represent a large fraction of the system cost.

**Submitted to WFCS'95  
IEEE International Workshop on  
Factory Communication Systems**

**October 4-6, 1995  
Ecole Polytechnique Fédérale de Lausanne  
Switzerland**

**Submission area: new protocols**

# **AbNET, a fault-tolerant fiber optic communication system**

**Harold Kirkham and Eddie Hsu  
Jet Propulsion Laboratory  
California Institute of Technology  
Pasadena, California, USA**

## **1 INTRODUCTION**

In this paper we will discuss a fiber optic communication system originally designed for monitoring and control of electric power distribution systems. The power system requirements lead to a novel design: the assumptions usually made in designing general-purpose computer communication systems do not apply. There are two particular differences. First, most local area networks (LANs) assume any user is likely to require communication with any other with equal probability, whereas the power system control application is hierarchical (master-slave), and peer-to-peer communication is rare. Second, LAN topologies are usually fixed in advance, and are rather simple (loops or buses). In contrast, the configuration of a monitoring system for power distribution may have to change in the short term because of damage, and in the long term because the power system evolves. There are advantages to having the system continue to operate even if some fibers are damaged, for system restoration.

The goal of the design effort was a communication system that could operate with any network topology, and would be reliable even if some fibers were damaged, or some nodes inoperative. The result was a set of protocols for operating a broadly applicable communication system, whose operation is transparent to the user. Because the flow of power in a distribution system is top-to-bottom, a hierarchical (master-slave) approach to control works well. However, peer-to-peer communications, for example to support MMS systems, could be performed with a modification to the original design.

## **2 DESIGN PHILOSOPHY**

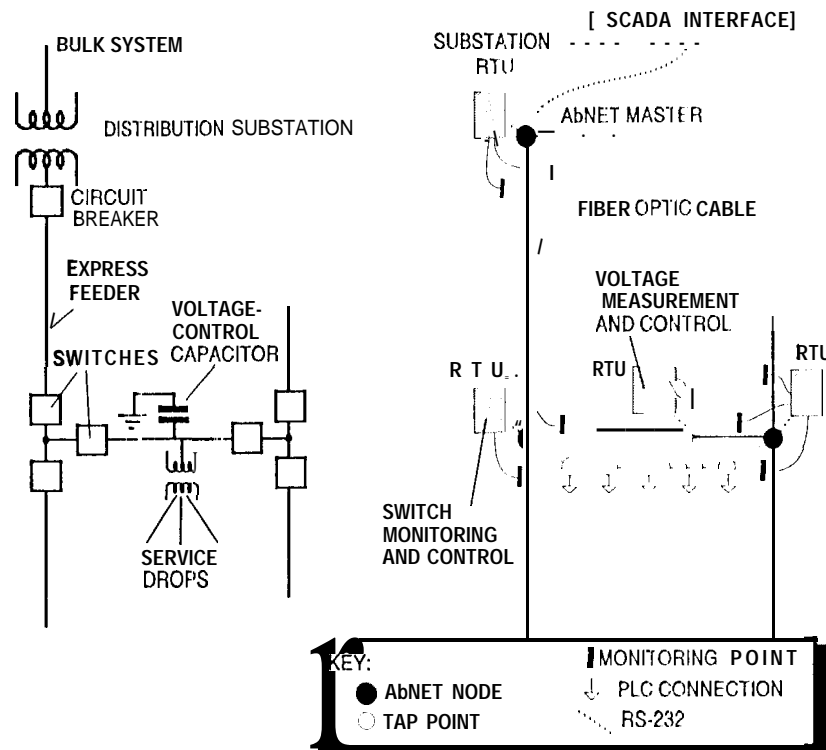
For control and monitoring of the power distribution network, a communication system that would meet the worst-case requirements was needed. This meant that

- it could access as many locations as necessary to support monitoring or control functions,
- it could handle the highest data rate likely to be required by any foreseeable application,
- and it would continue to operate even if part of the network were damaged.

Essentially, what was needed was a communication system that would be a transparent "phone company" for any and all automation functions. The topology of the communication network outside the substation (where control is assumed to originate), is congruent with the distribution system. This means that each time there is a lateral in the distribution system, for example, there must be a spur in the communication system. A master node would be located at each distribution substation. An example is shown in Figure 1.

## POWER SYSTEM


## COMMUNICATION SYSTEM



**Figure 1.** The communication system topology is fixed by the distribution system. Each communication node can handle branches in the fiber, and serves utility Remote Terminal Units

Normally, the distribution system is **operated** radially. There may be a handful of feeders from each substation, branching to serve the load. There is a limited set of loops that are normally open, but that can be reconfigured to provide an alternative way of bringing power to any given location. The fiber can, of course, cross a switch whether it is open or closed. As a result, the fiber optic communication system is arranged not as a conventional ring, star or bus system, but as a series of interconnected loops, with an occasional spur, as shown in Figure 2.

**Figure 2.** Typical AbNET system topology consists of multiple rings and loops



## 3 COMMUNICATION SYSTEM DESIGN

### 3.1 System Protocols

The problems to be solved in a communications network include message routing, access to the medium, error detection and congestion control. These problems have been addressed in a number of ways by network designers. A step towards standardization of computer

interconnections was taken when the Open Systems Interconnections model of the International Standards organization was released (Zimmerman, 1980).

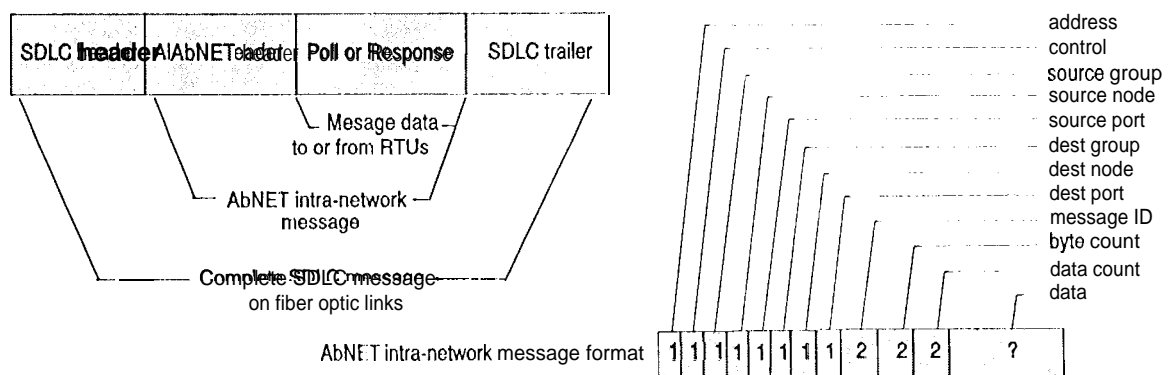
The model splits the problem of computer communications into seven logical and physical layers. While few existing networks adhere strictly to this breakdown (and some do not attempt to define any layer higher than the network layer), we have found it advantageous to place our consideration of the communications problem firmly within this framework.

### 3.2 Network Layer

We begin this description of the protocols with the network layer because it is this layer that most differentiates this system from others. The primary problems of the network layer are controlling the flow of data in the network, and routing within the network. None of the usual procedures adapt readily to the multiple-ring nature of our communications system. They do not solve the routing problem, nor can they easily prevent endless message circulation in a multiple-ring system. For routing our network uses a flooding algorithm. This is a distributed solution. At all nodes, any message received is retransmitted on all outgoing lines. By adopting such a strategy, a message inserted anywhere into the communications network will eventually be broadcast to all of the network. No message has any specific route; in fact all messages take all possible routes. This has two advantages: the shortest route is always taken, since all routes are always explored; and the transmission is very fault tolerant, since it does not depend on any particular path. The robust performance that this provides could be important in many industrial non-power applications.

The message must be removed, too, to prevent endless circulation. The header of every message contains a unique identifying number that can be stored in every node that it passes through. Each node decides whether or not the message will be repeated depending on whether the message has been seen before. The overhead required to implement this is small. One or two bytes, added to the message, will suffice for message numbering. The message format is shown in Figure 3.

This part of the communications protocol resembles the working of the body's immune system. On first exposure to a message, the nodes store information that will allow them, on a second exposure, to recognize and kill the message. It is because of this similarity that one of our group members proposed we call the system "AbNET," after the microbiologists' abbreviation "Ab," for antibody.



**Figure 3.** The AbNET message header contains group identification. It allows nodes to be selective about whose traffic they will handle. The AbNET package is embedded in a longer message, with its own header, for transmission on the fiber

The combined flooding/antibody method is quite efficient. The flooding algorithm means that all packets go down all links. The antibody algorithm means that all links transmit the message exactly once. Changes to the network configuration, even during operation, can easily be accommodated.

### 3.3 The Physical and Data Link Layers

The physical layer on the network side uses fiber optics as the transmission medium. Currently, the internodal speed is 2 Mbps, as speeds much higher than this would increase the cost. (Lower speeds, on the other hand, do not reduce the cost.) The data link layer uses SDLC type protocol. Because of the low bit error rate in the physical layer, we have not felt it necessary to implement error control, acknowledgment and retransmission in the data link layer. Any retransmission necessary due to error is handled by a layer above AbNET.

The network access between RTUs and AbNET is 1{ S-232 asynchronous. Use of RTS and CTS is made to decide on start and end of data frames. This allows total transparency of the protocol. This is important because several different RTU protocols are in use, sometimes even within one utility.

### 3.4 Medium Access

In previous power system communications, it has usually been necessary to limit traffic by having the RTUs make some decisions on their own. RTUs are often designed to operate with software that permits them to originate a transmission only if they detect some drastic change in the data they monitor. This approach is sometimes called Report-by-Exception. Some kind of medium access protocol must still be used. Collision detection has been used: the problem with this is that the response to a collision is always a delay. Information transfer is thus subject to unpredictable delays when a quick response is most needed.

Since the power system operating software is generally hierarchical, with centralized decision making and decentralized monitoring and control, The AbNET system uses a **centralized** polling strategy at the transport layer. An RTU is thus a "slave," and can transmit only if so directed by the master node. This method controls the slaves' access to the network and offers a mechanism for congestion control at a higher layer.

### 3.5 Zones

At the border between two service areas are gateway nodes. In most communication systems, a gateway is a node that can pass signals from the area of one communication system to another if required. Often, such a gateway will examine the destination address of a message, to see where to send it. In the AbNET system, such selectivity is reserved for the ordinary nodes. These nodes will only handle traffic associated with their master. The gateway nodes operate in what ethernet calls the "promiscuous" mode. They will handle any message presented to them. As a result of this approach, the nodes that are immediately adjacent to the border see additional traffic generated on the other side.

This reversal of normal roles has important advantages in the power system. First, at the border between two zones, or even two companies, only one communication node and one RTU need be installed, instead of two. Second, communications can keep pace with the changing power system configuration. Special messages can be sent from one zone into another to redefine ownership. Should power be lost to a substation unit, control of its service area can be recovered, and assigned to neighboring substations.

### 3.6 Services

A feature of a highly fault-tolerant network is that failure of any particular connection may not be detected, and several failures may have occurred before there is a communication failure. There are a number of ways this drawback can be overcome; one of the more obvious ones is to have intermediate nodes update a routing table in the message header. Changes in the route taken by messages can then be logged. This approach increases packet size, and may not detect changes in parts of the network that are redundant, because duplicate copies of packets are dropped.

An alternative is to temporarily suspend the distributed flooding algorithm, and instead use a part of the network layer that is designed specifically to have each node identify its nearest neighbor. The remote nodes build their own table of neighbors, which they send to the master periodically. Because the message flow required for this task resembles the operation of Sonar, we have designated this a Ping.

This approach takes time out of normal network operation. In the prototype network, a ping operation takes place about every 10 seconds, and a map at the master station is updated at this frequency.

### 3.7 System Speed

The requirements of the power system application dictated that the communication system be able to support data acquisition from a typical network at a rate of one complete scan every three seconds. We have estimated that this amounted to a data rate of about 3 kb/s on a feeder (Kirkham, Johnston and Allen, 1994), which on its face presents no challenge to a fiber based network.

However, while the fiber network may be underutilized at such a low data rate, the overall system may be challenged to meet the requirement. Suppose that a request for data is sent from the master station to the first node, and that the master station waits for a response before continuing. In the power system application, the RTU will be connected to the communications node via an RS-232 connection. If this is operated at 9600 b/s, and if the return data occupy 100 bytes, there will be a lag of about 100 ms before the polling sequence continues. This would limit the network size to approximately 30 RTUs. While this may be acceptable in many applications, it was thought that it should be possible to support a larger number. The polling approach was therefore modified. At present, four varieties of **polling** are available, as shown in Table 1.

Table 1. Comparison of protocol variants

Protocol type	Polling message		Response Message	
	Addressing mode	First hop relay method	First hop direction	Timing
AbNET 1	Individual	Flood	Flood	
AbNET 2	Broadcast	Onward only	Return path only	Simultaneous
AbNET 3	Broadcast	Flood	Flood	Repeat poll first
AbNET 4	Groups	Flood	Flood	Bursty

AbNET 1 is a polling approach in which the master waits for a response before moving on. In AbNET 2 and 3, the polling message is broadcast, and the slower RS-232 communications between the communication nodes and the RTUS take place at the same time throughout the network. (The difference between variants 2 and 3 is a small change in the sequence of events at the node after a poll command is received. AbNET 2 would immediately begin to pass the poll to the RTU, and repeat the poll message onward, whereas AbNET 3 would immediately flood the poll in all directions. In practice, the distinction may never be noticed.)

The AbNET 4 variant allows the master station to serve several groups of RTUS, and to distinguish between them. In this, it is the one variant that does not maintain complete transparency of use. It has the advantage that it can allow several "virtual networks" to co-exist on the same physical network, without interacting. This means that until there is complete interoperability of RTU protocols, RTUs with different protocols can safely be operated together. This variant was designed by Licom, the licensee of the AbNET system for power system applications.

With any of the variants of the original protocol, the time taken to poll the network for data is approximately the same as the time taken to poll a single RTU. While the amount of data in the system at any time is greater than in AbNET 1, it is still small by fiber standards. The maximum utilization factor is roughly the ratio of the RS-232 speed to the fiber speed, multiplied by the number of nodes. With 100 nodes running continuously, the fiber system would be loaded at about 50% capacity. With the normal 3-second scan, and with the anticipated amount of data, the fiber system would be loaded to less than 1% capacity. The spare capacity may find future use for other functions.

### 3.8 Reliability

The physics layer is assumed to have a bit error rate of less than  $10^{-9}$  for each link. With an average packet size of 100 bytes (approximately 1000 bits), about one packet in a million is corrupted and without error correction will be dropped by the receiving node. In a system of 100 nodes, scanned every 3 s, there are about  $10^9$  packets per year. Since the typical packet travels more than 1 hop, without redundancy more than 1000 packets per year will be dropped. (To avoid data loss, the SCADA system would re-poll if slaves fail to respond within some timeout period, essentially performing a transport layer function.)

For a packet transmission to fail in a typical AbNET network, it is necessary for transmissions on all links of a node to fail. Suppose there is a hypothetical network in which a node has 3 links, each connected to a master 10 hops away. (The typical degree of connectivity for each AbNET node is three. An AbNET node is rarely connected to more than three other nodes, and an AbNET node is not needed if it connects only two other nodes.) The probability of one of the 10 links *in all three paths* dropping the packet is  $(10 \times 10^{-6})^3$ , or  $10^{-15}$ , assuming there is no interconnection across paths. (Even this could be regarded as a worst-case scenario, as interconnection would improve the reliability.)

With flooding, it will be so long between dropped packets that other causes will predominate. The improvement due to the overconnected nature of the network is apparent, and justifies the non-use of error-detection-based retransmission at the data link layer.

While the numbers above are simple estimates based on assumed topologies, it is clear that the network is extremely robust if all the nodes are powered up. Power failure in part of the network is much more likely in the case of the power system application (where storm damage can black-out large areas), than in an industrial application, where the nodes are likely to be less remote. An industrial application can be expected to drop few packets.



#### 4 MMS

The AbNET protocols and the MMS protocols do not correspond well. Each set has communications following the flow of the product. In the AbNET case, the power flows from a central location out to the load; in the case of MMS, there is a "horizontal" flow corresponding to a production line. In the latter situation, peer-to-peer communication is normal, and valuable, whereas in the power system, top-to-bottom communication is more useful. One of the essential components of AbNET is *flooding*. While flooding actually ensures that all nodes receive all messages, it does not provide for peer-to-peer communication. This is because flooding is a dumb, distributed protocol, designed for robustness. While each node actually has available information about who the adjacent nodes are, no use is made of this, except at the master station. To abandon this philosophy would be to abandon the advantages of AbNET.

If peer-to-peer communication is demanded, it can be accomplished via the master, at the application layer. Bearing in mind that AbNET is ideally a transparent communication system, the application can use it to send peer-to-peer messages, with the master acting as a sort of data exchange. Such communication can only be initiated when the master polls the slave, so it may not be as fast as ordinary peer-to-peer communications. An estimate of speed can be made, however. The delay from peer to neighboring peer would be 4 RS-232 delays, instead of 2. This is hardly a serious impediment. An alternative scenario in which hierarchical and peer-to-peer communication co-exist, by a modification of the AbNET protocol, would require, in addition, the implementation of congestion control. This has not been required in any variant of AbNET.

In a scenario where there is a hierarchical relationship between a central command station and a collection of devices, the AbNET system would have an advantage over an architecture based on the peer-to-peer concept. Multiple devices may respond to a master simultaneously. If devices contain address filters of their own, then messages from the master can be directed toward a device, a group of devices or all of them. The flooding algorithm and the redundant topology ensure that all recipients can receive messages through the shortest path and with high reliability. Variants of AbNET provide a function resembling admission control, allowing the master to adjust the number of slaves participating in the hierarchical conversation with the master. In power system applications, this number is likely to be fixed. However, a dynamic algorithm for determining the number may be implemented to control, in addition to the number of master-slave conversations, the number of peer-to-peer conversations allowed in the network at any given time.

#### 5 CONCLUSION

The design of a fiber optic communication network for monitoring and control in power distribution systems has been discussed. By appropriate choice of protocols, a fault-tolerant system that operates in any arbitrary network configuration has been devised.

The network, called AbNET, is a packet-based distributed protocol system. Flooding is used for maximum failure tolerance. Hierarchical (master-slave) polling controls access to the system. This supports many data acquisition and control applications. The protocols allow multiple adjacent masters to share resources. A low-level network service reports the network's configuration to the master, where changes can be logged, and action taken if needed.

The system is transparent to the user, and maintains no record of the clients it serves. It is fast enough for many industrial control applications. Because hierarchical access control is used, peer-to-peer communications, such as for MMS systems, do not map well into the

protocols, but can be accomplished via software at the master station.

Only a small number of fiber cables would be needed for a high reliability system. In many industrial applications, where the inter-node distance is not large, the fiber is unlikely to represent a large fraction of the system cost.

The AbNET system offers a reliable architecture through the use of fiber optics, a flooding algorithm and redundant connections. The hierarchical property offers an immediate alternative in connecting factory systems that exhibit hierarchic] characteristics. In addition, peer-to-peer communications may be accommodated through the use of the master node. We encourage those who are interested to consider the use of AbNET in other areas.

## 6 REFERENCES

- Zimmermann, H., (1980) *OSI Reference Model - the ISO Model of Architecture for Open Systems Interconnection* IEEE Transactions on Communications, Volume COM-28, No. 4, April, pp. 425-432.
- Kirkham, H., Johnston, A. R. and Allen, G. D. *Design considerations for a fiber optic communications net work for power systems*, IEEE Trans. Power Delivery, Vol 9, No 1, Jan 1994, pp 510-518.

## 7 ACKNOWLEDGEMENTS

The authors would like to acknowledge the contributions of Shannon Jackson and Maclen Marvit in the design and testing of network prototype hardware, and in writing simulation software and GUI software.

Prepared by the Jet Propulsion Laboratory, California Institute of Technology, for the U.S. Department of Energy, Office of Energy Management Systems, Utility Systems Division, through an agreement with the National Aeronautics and Space Administration.

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, product, or process disclosed, or represents that its use would not infringe privately owned rights.